

TROPICAL GEOMETRY AND MECHANISM DESIGN

ROBERT ALEXANDER CROWELL AND NGOC MAI TRAN

ABSTRACT. We develop a novel framework to construct and analyze finite valued mechanisms using tropical convex geometry. Our main results geometrically characterize all possible incentive compatible mechanisms together with their payments on arbitrary type spaces. We obtain a geometric interpretation of revenue equivalence which allows for a more refined analysis than previous results. A distinct feature of our theory is that it is constructive.

1. INTRODUCTION

Mechanisms are engineered games, devised to implement outcomes that depend on the private information of individuals in the economy. They provide a theoretical model to study which economic allocations an institution can achieve as equilibrium of a game in which individuals' preferences, called their *type* is not known by the institution. When this information is relevant for taking a decision, and individuals can misreport their preferences in order to manipulate the outcome in their favor, they need to be incentivized to reveal it truthfully. A mechanism that elicits agents' types truthfully is said to be *incentive compatible* (IC).

In this paper we focus on single agent mechanisms, which form a fundamental building block of dominant strategy mechanism design, *cf.* [11, 27]. The analysis of such mechanisms is predominantly analytic or algebraic [20, 22, 27]. We propose an entirely different approach via tropical convex geometry and tropical combinatorics. Within this framework, the study of incentive compatibility becomes a question about tropical eigenspaces and point configurations in the tropical affine space \mathbb{TP} . To handle these economic problems geometrically, we put forward the definition of the basic region $\mathbf{basic}(T)$ and the set of basic cells $\mathbf{cells}(T)$ for an arbitrary multiset of types $T \subset \mathbb{TP}^{m-1}$. Our first main result states that any IC mechanism on T and its payments can be constructed from these objects.

Theorem 1.1. *Let $T \subset \mathbb{TP}^{m-1}$ be a multiset of types, and (g, p) a single agent mechanism with outcome function $g : T \rightarrow [m]$ and payment vector $p \in \mathbb{TP}^{m-1}$.*

- (1) *A mechanism (g, p) is IC if and only if p lies in a basic cell $P \in \mathbf{cells}(T)$, and g is contained in $\mathbf{coVec}_T(P)$ as a subgraph. Then P is the set of all IC payments of g .*
- (2) *A vector $p \in \mathbb{TP}^{m-1}$ is the payment function of some IC mechanism (g, p) if and only if p lies in the basic region $\mathbf{basic}(T)$.*
- (3) *If T is finite and generic, then the basic cells $\mathbf{cells}(T)$ are full-dimensional cells of the max-plus tropical polytope generated by T .*
- (4) *Let $(T^n, n \geq 1)$ be a sequence of finite sets which approximate T . Then the basic cells in $\mathbf{cells}(T)$ are limits in the Hausdorff metric of the basic cells of T^n as $n \rightarrow \infty$.*

Theorem 1.1 supplies a novel framework to study IC mechanisms, and more importantly, the set of possible IC payments, on a given type space. In particular, parts (1) and (2)

2010 *Mathematics Subject Classification.* 91A80, 14T05.

Key words and phrases. Mechanism Design, Tropical Geometry, Tropical Convexity, Tropical Combinatorics, Incentive Compatibility, Cyclical Monotonicity, Revenue Equivalence.

state that the basic cells $\mathbf{basic}(T)$ and their covectors determine the set of all IC mechanisms and IC payments on T , while parts (3) and (4) say that they are easy to discern from the geometry of T .

Our second main result, Theorem 4.3, gives a geometric construction and meaning to important classes of mechanisms, such as weakly monotone and revenue equivalent mechanisms on an arbitrary type space T . This theorem relates geometry and linear algebra, giving a clear understanding of how the outcome function interacts with the geometry of the type space in affecting these properties.

The geometric approach developed here is conceptually distinct from large parts of the existing literature on mechanism design. Our theorems provide new tools that make checking and visualizing incentive compatibility easy, both theoretically and computationally. Furthermore, Theorem 1.1 and 4.3 provide powerful techniques to *construct* all possible IC outcome functions together with their payments for arbitrary type spaces. To our knowledge such a joint characterization has not been obtained previously and appears to be beyond the scope of the traditional mathematical machinery of mechanism design.

Let us demonstrate the insights that can be gained using these theorems. An IC mechanism is *revenue equivalent* (RE) if there exists a unique IC payment for its outcome function, up to an additive constant. We shall say that a type space T is revenue equivalent if all IC mechanisms on T are revenue equivalent. Prior to our work, the characterization of RE type spaces, due to Chung and Olszewski [6], reads as follows: T is RE if and only if there do not exist disjoint subsets B_1, B_2 , a function $r : B_1 \cup B_2 \rightarrow \mathbb{R}$, and an $\epsilon > 0$, such that T equals the union of two non-empty sets $V_+(B_1, \epsilon, r)$ and $V_-(B_2, \epsilon, r)$, whose definitions depend on the parameters given. Given a specific type space T , it is not immediate how to apply this theorem to check whether T is RE. In contrast, we obtain the following characterization as a consequence of Theorem 1.1 and 4.3.

Theorem 1.2. *A type space $T \subset \mathbb{TP}^{m-1}$ is RE if and only if for each $p \in \mathbf{basic}(T)$ the graph of p is strongly connected.*

The graph of p is the directed graph on n nodes, with edge (i, j) if and only if the distance between $T \cap \overline{\mathcal{H}}_i(-p)$ and $\overline{\mathcal{H}}_{ij}(-p)$ is zero, where $\overline{\mathcal{H}}_i(-p)$ and $\overline{\mathcal{H}}_{ij}(-p)$ denote max-plus half-spaces of type $\{i\}$ and $\{i, j\}$ respectively. The max-plus hyperplane $\overline{\mathcal{H}}(-p)$ at p equals the max-plus hyperplane at 0 translated by p , so that checking whether a type space T is RE using Theorem 1.2 only requires considering translations of a fixed set by points in $\mathbf{basic}(T)$ and recording which sectors have positive distances. This makes checking revenue equivalence for a given type space surprisingly simple, while such an algorithm is not apparent from Chung and Olszewski's characterization. In addition, Theorem 1.1 and 4.3 allow for more detailed conclusions. For example, on an arbitrary type space T , one can precisely identify the set of all RE and non-RE mechanisms. We illustrate these techniques in Examples 4.6 and 4.7.

To demonstrate the insights gained from the combinatorial view, consider the following problem. Suppose $T \subset \mathbb{TP}^{m-1}$ is a multiset consisting of r points. Let $d(T) \in \mathbb{N}$ be the number of IC outcome functions on T . We are not aware of ways to compute or bound $d(T)$ using traditional techniques. On the other hand, this easily follows from Theorem 1.1.

Corollary 1.3. *Suppose $T \subset \mathbb{TP}^{m-1}$ consists of r points. Then $1 \leq d(T) \leq \binom{r-1}{m-1}$. If T is generic, then $d(T) = \binom{r-1}{m-1}$.*

Our theorems were obtained by applying results from tropical convexity and tropical combinatorics to problems in mechanism design. However, we go further, by adapting and developing tropical convex geometry to handle economic problems. The basic region $\text{basic}(T)$ and basic cells $\text{cells}(T)$ of a multiset T in tropical affine space are entirely new objects in tropical convex geometry, motivated by incentive compatibility problems. Thus concepts in mechanism design lend useful interpretations to objects in tropical convex geometry. Conversely the mathematical apparatus we advocate in this paper allows for the analysis of mechanisms in a unified way. It applies equally to mechanism design with and without money, and to multidimensional mechanism design. At the same time it supplies a complementary geometric view to the existing edifice of mechanism design theory, allowing for additional insights, as demonstrated in the preceding paragraphs.

Mechanism design is a cornerstone of economic theory with many prominent applications, ranging from the analysis and design of voting procedures [13, 19], trading and auctioning mechanisms [20, 21, 26], to contract design and regulatory policy [17]. The plenitude of open problems, both theoretical and applied, make it a very active research area. Tropical mathematics offers a new tool box to study and interpret mechanisms, with a novel perspective on open problems. This promises further developments between tropical geometry and mechanism design, though we certainly do not claim that all problems in mechanism design can be solved tropically. More broadly, there have been applications of tropical geometry to other economic problems, such as mean-payoff games [1], trade theory [14, 24], and auction theory [3, 25]. These papers attest to a fruitful interaction between tropical geometry and economics.

Further Literature. This paper is written for a general math audience with minimal background in either mechanism design or tropical geometry. While the paper is self-contained, it is not intended as an introduction to either field. For streamlined exposition with minimal notation, we only consider the simplest case of dominant strategy mechanism design, namely, direct mechanisms with a finite set of outcomes and one agent in a quasi-linear setting. These can be embedded into richer models. For an in-depth treatment of mechanism design, with economic interpretations and further generalizations, we refer to the excellent monograph of Vohra [27]. Similarly, for a fuller picture of tropical algebra and geometry, there exists the comprehensive texts of Butkovic [5] and Bacelli et al. [2] on solving tropical linear equations, or Joswig [15] and Maclagan and Sturmfels [18] on tropical geometry.

Organization. We introduce the essential background in mechanism design and tropical convex geometry in Section 2. In Section 3, we prove Theorem 1.1 by developing auxiliary results along the way. In particular we formalize the concept of basic cells and the basic set. In Section 4, we state and prove Theorem 4.3, providing a series of examples to demonstrate its power, and including a proof of Theorem 1.2 stated above.

Notation. For an integer n , define $[n] := \{1, 2, \dots, n\}$. For sets $A, B \subset \mathbb{R}^m$, write $\text{d}(A, B)$ for the infimum of the distance between pairs of points in these sets, and $\text{d}_H(A, B)$ for their Hausdorff distance, with respect to the Euclidean norm. We shall use the underline notation, such as $\underline{\oplus}, \underline{\odot}, \underline{\mathcal{H}}, \dots$ to indicate objects defined with arithmetic done in the min-plus algebra, and the overline notation $\overline{\oplus}, \overline{\odot}, \overline{\mathcal{H}}, \dots$ to indicate the same objects defined with arithmetic in the max-plus tropical algebra. When we write a multiset, we will use the same notation as for sets. When listing or enumerating the elements, we will list copies. The cardinalities of multisets are understood to include copies. Identify a graph with its incidence matrix.

2. MECHANISM DESIGN, TROPICAL ALGEBRA AND TROPICAL CONVEX GEOMETRY

2.1. Mechanism Design in Tropical Algebra. Consider a game with one agent and $m \in \mathbb{N}$ possible outcomes. Fix $T \subset \mathbb{R}^m$, called the *type space*. At the beginning of the game, Nature chooses a true type $t^* \in T$ which only the agent knows. The i -th coordinate t_i^* measures how much the agent values outcome i . A *mechanism* is a pair (g, p) , consisting of an outcome function $g : T \rightarrow [m]$, which is always assumed to be onto, and a payment function $p : T \rightarrow \mathbb{R}$. The agent's action is to declare to the mechanism a type $s \in T$, which may be different from the true type t^* . If s is declared, the game's outcome is $g(s)$, and the agent needs to pay $p(s)$. In this case, the agent's utility is

$$u([g(s), p(s)], t^*) = t_{g(s)}^* - p(s).$$

The agent, knowing (g, p) , will declare a type $s \in T$ that maximizes utility. A central goal of mechanism design is to identify *incentive compatible mechanisms*, these are mechanisms under which the agent will always tell the truth.

Definition 2.1. Say that a mechanism (g, p) is *incentive compatible* (IC) if regardless of t^* , the agent always maximizes utility by declaring the true type. That is,

$$(IC) \quad t_{g(t^*)}^* - p(t^*) \geq t_{g(s)}^* - p(s) \quad \text{for all } s, t^* \in T.$$

Say that an outcome function $g : T \rightarrow [m]$ is incentive compatible if there exists $p : T \rightarrow \mathbb{R}$ such that (g, p) is IC.

We now proceed to state the essential terminology of mechanism design using tropical mathematics. Tropical linear algebra and its geometric version, tropical convex geometry, is the study of matrices, linear spaces and convex sets with arithmetic done in the tropical semi-ring. The *min-plus semi-ring* $(\mathbb{R} \cup \{+\infty\}, \underline{\oplus}, \odot)$ has addition and multiplication defined by

$$a \underline{\oplus} b := \min(a, b), \quad a \odot b := a + b \quad \text{for } a, b \in \mathbb{R} \cup \{+\infty\}.$$

Analogously, the *max-plus semi-ring* $(\mathbb{R} \cup \{-\infty\}, \overline{\oplus}, \odot)$ is defined by

$$a \overline{\oplus} b := \max(a, b), \quad a \odot b := a + b \quad \text{for } a, b \in \mathbb{R} \cup \{-\infty\}.$$

A matrix $L \in \mathbb{R}^{m \times m}$ is said to have *min-plus eigenvalue-eigenvector* pair $(\lambda, p) \in \mathbb{R} \times \mathbb{R}^m$ if

$$L \odot p = \lambda \odot p,$$

where the matrix-vector and scalar-vector multiplications take place in the min-plus semiring. Explicitly, for all $i = 1, \dots, m$,

$$\min_{j=1, \dots, m} L_{ij} + p_j = \lambda + p_i.$$

By a theorem of Cuninghame-Green [8], a $m \times m$ matrix with finite entries $L \in \mathbb{R}^{m \times m}$ has a unique min-plus eigenvalue, interpreted as the smallest normalized cycle length on a graph with edge weights given by L . Thus one can speak of the *min-plus eigenvalue* of a matrix L , denoted $\underline{\lambda}(L)$. The *min-plus eigenspace* $\underline{\text{Eig}}(L)$ of $L \in \mathbb{R}^{m \times m}$ is its set of eigenvectors

$$\underline{\text{Eig}}(L) = \{x \in \mathbb{R}^m : L \odot x = \underline{\lambda}(L) \odot x\}.$$

Returning to mechanism design, the *allocation matrix* L^g of a given outcome function $g : T \rightarrow [m]$ is an $m \times m$ matrix, with jk -th entry

$$(1) \quad L_{jk}^g = \inf_{t \in g^{-1}(j)} \{t_j - t_k\}.$$

As is common in the literature [27], we shall always assume T and g to be such that the entries of L^g are finite. It follows from Definition 2.1 that $p(t) = p(s)$ whenever $g(t) = g(s)$, for any IC mechanism (g, p) . Thus we may assume that $p \in \mathbb{R}^m$. By algebraic manipulations, one can check that equation (IC) holds if and only if the matrix $L^g \in \mathbb{R}^{m \times m}$ has a tropical eigenvector $p \in \mathbb{R}^m$ with tropical eigenvalue zero. That is, (g, p) is an IC mechanism if and only if

$$(2) \quad L^g \odot p = 0 \odot p = p.$$

Hence, deciding IC for a given g involves solving the tropical linear equation (2). Cuninghame-Green's theorem applied to equation (2) gives the classical characterization of IC outcome functions obtained by Rochet [22].

Theorem 2.2 ([8], [22]). *An outcome function g is IC if and only if the matrix L^g has min-plus tropical eigenvalue zero, that is, $\underline{\lambda}(L^g) = 0$.*

Definition 2.3. If an outcome function $g : T \rightarrow [m]$ is IC, call $\underline{\text{Eig}}(L^g)$ the set of *incentive compatible payments* of g .

Definition 2.4. Say that g is *weakly monotone* if the matrix $L^g + (L^g)^\top$ is element-wise nonnegative.

Definition 2.5. Say that g is *revenue equivalent* (RE) if it is IC, and its set of incentive compatible payments $\underline{\text{Eig}}(L^g)$ consists of exactly one point up to tropical scalar multiplication. That is, g is RE if for any pair $p, q \in \underline{\text{Eig}}(L^g)$, there exists a constant c such that $p_i = q_i + c$ for all $i = 1, \dots, m$. Say that a type space is *revenue equivalent* if any incentive compatible mechanism defined on it is revenue equivalent.

2.2. Mechanism design and tropical convex geometry. Our key departure from the network flow approach to mechanism design, developed in [27], will be the use of tropical convex geometry. This subsection shall make Theorem 1.1 precise. In a first step we replace type sets in \mathbb{R}^m by multisets in tropical affine space \mathbb{TP}^{m-1} . In a second step we define fundamental concepts in tropical convex geometry including tropical hyperplanes, polytopes, covectors, basic cells and basic sets.

The $(m - 1)$ -dimensional *tropical affine space* \mathbb{TP}^{m-1} is \mathbb{R}^m modulo the line spanned by the all-one vector

$$\mathbb{TP}^{m-1} \equiv \mathbb{R}^m / \mathbb{R} \cdot (1, \dots, 1).$$

We denote by $\pi : \mathbb{R}^m \rightarrow \mathbb{TP}^{m-1}$ the canonical projection from \mathbb{R}^m to \mathbb{TP}^{m-1} . The space \mathbb{TP}^{m-1} arises when we want to regard a tropical vector $t \in \mathbb{R}^m$ and its scalar multiple $a \odot t = (a + t_1, \dots, a + t_m)$ as equivalent. Whenever convenient, such as visualizing examples, we shall identify \mathbb{TP}^{m-1} with \mathbb{R}^{m-1} via the homeomorphism

$$(3) \quad \{a \odot (x_1, \dots, x_m) : a \in \mathbb{R}\} \in \mathbb{TP}^{m-1} \mapsto (x_2 - x_1, \dots, x_m - x_1) \in \mathbb{R}^{m-1}.$$

A *multiset* in \mathbb{TP}^{m-1} is a set of unique elements in \mathbb{TP}^{m-1} together with a function counting the number of copies of the underlying set's elements. A multiset in which each element has only one copy is a set. As a convention, we shall write a multiset by listing copies of its unique elements. Functions defined on a multiset may assign distinct values to the copies of the unique elements.

To translate incentive compatibility problems to \mathbb{TP}^{m-1} , we proceed as follows. Given a type set $T' \subset \mathbb{R}^m$, define the multiset $T \subset \mathbb{TP}^{m-1}$ with underlying set $T = \pi(T')$, by

assigning to each $t \in T$, m copies whenever $\pi^{-1}(t) \cap T'$ contains at least m points, and k copies whenever $\pi^{-1}(t) \cap T'$ contains $k < m$ points. Note that this definition depends only upon the type set $T' \subset \mathbb{R}^m$. Suppose now we were given an outcome function $g' : T' \rightarrow [m]$. We can then define an induced outcome function $g : T \rightarrow [m]$ which maps the copies of $t \in T$ to the values of g' on $\pi^{-1}(t) \cap T'$. Note that $L^g = L^{g'}$, so that g is IC if and only if g' is. Furthermore, if p is an IC payment of g , then so is $a \odot p$ for all $a \in \mathbb{R}$, so that IC payments naturally are subsets of \mathbb{TP}^{m-1} . Hence, for incentive compatibility problems, there is no loss in passing from g' , an outcome function defined on a subset of \mathbb{R}^m , to g , an outcome function defined on a sufficiently rich multiset in \mathbb{TP}^{m-1} . In economic terms, this means only the relative valuation of the agent matter for truthfulness.

In the following we recall some fundamental concepts in tropical convex geometry. Most definitions we build on appeared in the seminal work of Develin and Sturmfels [9], and were developed further in subsequent works by the authors of [10, 12, 16]. However, tropical convex geometry itself has a much longer history, for a comprehensive introduction consult [15, 18] and references therein.

For a point $t \in \mathbb{TP}^{m-1}$, the *min-plus hyperplane with apex t* , denoted $\underline{\mathcal{H}}(-t)$, is the set of $z \in \mathbb{TP}^{m-1}$ such that the minimum in the tropical inner product

$$(4) \quad (-t)^\top \odot z = \min\{z_1 - t_1, \dots, z_m - t_m\}$$

is achieved at least twice. For a subset $I \subseteq [m]$, denote by $\underline{\mathcal{H}}_I(-t)$ the *min-plus half-space I* of $\underline{\mathcal{H}}(-t)$, which is the set of $z \in \mathbb{TP}^{m-1}$ such that the minimum in the tropical inner product is achieved at all indices $i \in I$ (and possibly more). That is,

$$\underline{\mathcal{H}}_I(-t) = \{z \in \mathbb{TP}^{m-1} : z_i - t_i \leq z_j - t_j \text{ for all } i \in I, j \in [m]\}.$$

When $I = \{i\}$ or $I = \{i, j\}$, we write $\underline{\mathcal{H}}_i$ and $\underline{\mathcal{H}}_{ij}$, respectively, instead of $\underline{\mathcal{H}}_I$. We shall call $\underline{\mathcal{H}}_i$ the *i -th sector* of the hyperplane.

For a multiset $T \subset \mathbb{TP}^{m-1}$, the union of min-plus hyperplanes $\bigcup_{t \in T} \underline{\mathcal{H}}(-t)$ partitions \mathbb{TP}^{m-1} into a polyhedral complex, called the *min-plus hyperplane arrangement* based on T , denoted $\underline{\mathcal{H}}(-T)$. For a finite set T , the set of bounded cells of $\underline{\mathcal{H}}(-T)$ is the *max-plus polytope* generated by T [9, Theorem 15]. Say that a multiset $T \subset \mathbb{TP}^{m-1}$ is *generic* if there is no subset of $2 \leq k \leq m$ points in T whose projection onto k coordinates lie on a tropical hyperplane in \mathbb{TP}^{k-1} . Note that a generic multiset T necessarily has no copies, i.e. it is a set.

The *min-plus covector* of a point $p \in \mathbb{TP}^{m-1}$ with respect to a multiset $T \subseteq \mathbb{TP}^{m-1}$, denoted $\text{coVec}_T(p)$, is the bipartite graph with nodes $[m] \times T$ and edge $(i, t) \in [m] \times T$ if and only if $p \in \underline{\mathcal{H}}_i(-t)$.

Definition 2.6. Fix a multiset $T \subset \mathbb{TP}^{m-1}$. Let g be a bipartite graph on $[m] \times T$. We will say that g is an *outcome graph* if it defines an outcome function $g : T \rightarrow [m]$ via $g(t) = i$ if and only if (i, t) is an edge of the graph g .

Definition 2.7. To a bipartite graph g on $[m] \times T$, the possibly empty *polyhedron of g* is

$$(5) \quad \mathcal{P}(g) = \{q \in \mathbb{TP}^{m-1} : q \in \underline{\mathcal{H}}_i(-t) \text{ for all } (i, t) \in g\}.$$

By min-max duality, $\mathcal{P}(g)$ can be equivalently written as

$$(6) \quad \mathcal{P}(g) = \{q \in \mathbb{TP}^{m-1} : t \in \overline{\mathcal{H}}_i(-q) \text{ for all } (i, t) \in g\}.$$

It is immediate that if a polyhedron $\mathcal{P}(g)$ is non-empty, then all points in its relative interior have the same covector with respect to T . Call this the *basic covector* of g , denoted $\nu(g)$.

Definition 2.8. Say that a covector ν is basic, and that $\mathcal{P}(\nu)$ is a *basic cell*, if $\nu = \nu(g)$ for some outcome function g . The set of basic cells of T is

$$\text{cells}(T) = \{\mathcal{P}(\nu) : \nu \text{ is a basic covector}\}.$$



The set union of the basic cells, denoted

$$\text{basic}(T) = \bigcup_{P \in \text{cells}(T)} P.$$


will be called the *basic set of T* .

Example 2.9 (Hyperplanes, half-spaces and basic cells). Panel (A) in Figure 1 depicts the min-plus hyperplane $\underline{H}(-t_0)$ and its sectors (top) and the max-plus hyperplane $\overline{H}(-t_0)$ and its sectors (bottom). Both are in \mathbb{TP}^2 with apex t_0 . As usual, we identify \mathbb{TP}^2 with \mathbb{R}^2 via the map given in (3).

Example 2.10 (Covectors and basic cells). Panel (B) depicts a set $T = \{t_1, t_2, t_3\}$, the min-plus arrangement $\underline{H}(-T)$, and the covectors of a selection of cells. There is only one basic cell, shaded gray. Its covector defines the outcome function $g : T \rightarrow [3]$ with $g(t_1) = 1, g(t_2) = 2, g(t_3) = 3$. By Theorem 1.1, g is the only IC outcome function amongst all possible outcome functions on T , and its set of IC payments is the basic cell shaded gray. One can also apply Corollary 1.3: T consists of three generic points, so the number of IC outcome functions on T is $\binom{m-1}{m-1} = 1$.



(A) Tropical min-plus and max-plus hyperplanes



(B) A min-plus arrangement on three points labeled with some of its covectors and a single basic cell shaded gray.

FIGURE 1. Hyperplanes and arrangements. Figures accompany Examples 2.9 and 2.10. For all figures in this paper, axis orientation and sector labels follow the convention set in Panel (A).

3. AUXILIARY RESULTS AND PROOF OF THE MAIN THEOREM

In this section, we prove Theorem 1.1 by establishing a series of lemmata and propositions. For readability, we split the Theorem into three propositions followed by their proofs. Proposition 3.3 covers parts (1) and (2) of Theorem 1.1. Parts (3) and (4) are covered by Propositions 3.7 and 3.11, respectively. In the process we clarify all subtleties and discuss some examples to aide understanding.

3.1. Covectors and incentive compatibility. An outcome function $g : T \rightarrow [m]$ is onto, and each $t \in T$ must be assigned to an outcome. Thus, a bipartite graph g is an outcome graph if and only if each node $t \in T$ has degree one, and each node $i \in [m]$ has degree at least one. The goal is to discern which outcome functions g are IC. Proposition 3.3 gives a characterization in terms of the covector of the polyhedron $\mathcal{P}(g)$.

Lemma 3.1. *Let $g : T \rightarrow [m]$ be an outcome function. Then $\mathcal{P}(g)$ is the set of incentive compatible payments of g . In particular, $\mathcal{P}(g) \neq \emptyset$ if and only if g is IC.*

Proof. Note that $p \in \mathcal{P}(g)$ if and only if for all $t \in g^{-1}(i) \subset T$ and all $j \in [m]$, $p_i - t_i \geq p_j - t_j$. Thus p is an incentive compatible payment of g . \square

Lemma 3.2. *Let h be a bipartite graph on $[m] \times T$. The following are equivalent.*

- (1) $\mathcal{P}(h)$ is not empty,
- (2) h is contained as a subgraph in the covector of some cell of $\underline{\mathcal{H}}(-T)$,
- (3) $\mathcal{P}(h) = \mathcal{P}(\nu(h))$ for the covector $\nu(h)$ of a point in the relative interior of $\mathcal{P}(h)$.

Proof. Suppose (1). Let p be a point in the relative interior. By definition of covector, all points in the relative interior of $\mathcal{P}(h)$ have the same covector. Let us denote this covector by $\nu(h)$. If $(i, t) \in h$, then $p \in \underline{\mathcal{H}}_i(-t)$, so $(i, t) \in \nu(h)$. Thus h is contained in $\nu(h)$. This implies (2). In addition, it also implies $\mathcal{P}(\nu(h)) \subseteq \mathcal{P}(h)$, by definition of the polyhedra $\mathcal{P}(h)$ and $\mathcal{P}(\nu(h))$. Now, $\mathcal{P}(h)$ and $\mathcal{P}(\nu(h))$ are closed polyhedra. So if $\mathcal{P}(\nu(h))$ is a strict subset of $\mathcal{P}(h)$, then there must exist points in the relative interior of $\mathcal{P}(h)$ which do not belong to $\mathcal{P}(\nu(h))$, and in particular, cannot have covector $\nu(h)$. This is a contradiction, so $\mathcal{P}(\nu(h)) = \mathcal{P}(h)$. This establishes (3). Finally, (3) trivially implies (1). \square

Proposition 3.3 (Covector characterization). *Let $g : T \rightarrow [m]$ be an outcome function on a multiset T . Then g is IC with payment $p \in \mathbb{TP}^{m-1}$ if and only if $\underline{\text{coVec}}_T(p)$ contains g as a subgraph. In this case, $\mathcal{P}(\nu(g))$ is the set of incentive compatible payments of g .*

Proof. By Lemma 3.1 the mechanism (g, p) is IC if and only if $p \in \mathcal{P}(g)$. By Lemma 3.2 it is without loss to assume that p is in the relative interior, with covector $\underline{\text{coVec}}_T(p)$. It follows from the same lemma that $p \in \mathcal{P}(g)$ if and only if g is contained in $\underline{\text{coVec}}_T(p)$ and $\mathcal{P}(\underline{\text{coVec}}_T(p)) = \mathcal{P}(g)$. \square

In other words, to each IC outcome function there corresponds a basic cell containing its payments. Conversely, each basic cell can be associated with a set of IC outcome functions whose payments it contains. This set consists of precisely those outcome functions contained in the cell's covector.


Example 3.4 (Algebraic construction of IC mechanisms). Let ν_1 (resp. ν_2) be the cells containing p (resp. q) in its relative interior. The covectors of p and q are given by

$$\underline{\text{coVec}}_T(p) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \underline{\text{coVec}}_T(q) = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}.$$

There is a unique outcome function $g : T \rightarrow [3]$ whose set of IC payments equals ν_1 , namely, $g(t_1) = 1, g(t_2) = 2, g(t_3) = 3$. For the cell ν_2 , there is no outcome function contained in $\underline{\text{coVec}}_T(q)$, so there exists no onto IC mechanism with payment set ν_2 .

Example 3.5 (Geometric construction of IC mechanisms). In Figure 2, consider the max-plus hyperplane with apex p (dotted). The sets $\overline{\mathcal{H}}_i(-p) \cap T$ for $i \in [m]$ partition T into non-empty, disjoint subsets. By Lemma 3.1 and (6), any outcome function g such that (g, p) is IC must map types in the interior of $\overline{\mathcal{H}}_i(-p) \cap T$ to i , for $i \in [3]$. Thus $g(t_1) = 1, g(t_2) = 2, g(t_3) = 3$ is the unique such outcome function, agreeing with Example 3.4.

Example 3.6 (Non-basic cells cannot contain IC payments). In Figure 2 (B), consider the max-plus hyperplane with apex q (dotted). Suppose for contradiction that there exists an outcome function g such that (g, q) is IC. By (6) and Lemma 3.1, $g(t_1) = g(t_3) = 2$, while $t_2 \in \overline{\mathcal{H}}_{13}(-q) \cap T$, $g(t_2)$ must be either 1 or 3. But in either assignment of t_2 , g is not onto, so it cannot be an outcome function. Thus, q cannot be an IC payment of some IC mechanism.



(A) Construction of an IC mechanism by assigning types in each sector of a max-plus tropical hyperplane with apex p to the corresponding outcomes.

(B) No IC outcome function for non-basic cells. Defining an onto outcome function is not possible because there are too few types in some sectors.

FIGURE 2. Here $T = \{t_1, t_2, t_3\}$ with no repeated points. The min-plus arrangement $\underline{\mathcal{H}}(-T)$ is drawn using solid lines. Axis orientation and sector labels follow the convention set in Figure 1. The unique basic cell is shaded gray. Figure accompanies Examples 3.4 to 3.6.

3.2. Basic cells via generic approximations. By Proposition 3.3, the set of basic cells $\text{cells}(T)$ and its union, the basic set $\text{basic}(T)$, encode the IC payments of any possible mechanism. We now study some properties of basic cells, starting with the case where T is a multiset of generic points. Recall that this means there does not exist a subset of $2 \leq k \leq m$ points in T whose projection onto k coordinates lie on a tropical hyperplane in \mathbb{TP}^{k-1} . Thus, in particular, T is a *set*.

Proposition 3.7. *Suppose $T \subset \mathbb{TP}^{m-1}$ contains r generic points. Then $\text{cells}(T)$ is precisely the set of full-dimensional cells of $\underline{\mathcal{H}}(-T)$. In particular, the cardinality of $\text{cells}(T)$ is $\binom{r-1}{m-1}$.*

Proof. By [9, Corollary 12], $P = \mathcal{P}(\nu)$ is a full-dimensional cell of $\underline{\mathcal{H}}(-T)$ if and only if ν is an outcome graph whose polyhedron $\mathcal{P}(\nu)$ is non-empty. Thus, the set of full-dimensional cells of $\underline{\mathcal{H}}(-T)$ is a subset of $\text{cells}(T)$. On the other hand, let $P = \mathcal{P}(\nu) \in \text{basic}(T)$ be a basic cell with respect to some outcome graph g . Since ν is basic, $\nu = \nu(g) = g$, so P is a full-dimensional cell of $\underline{\mathcal{H}}(-T)$. Finally, the cardinality follows from [9, Corollary 25]. \square

To handle the case of a general multiset T we use approximations by generic perturbations, which are common technique in tropical geometry, see [18].

Definition 3.8. Suppose the multiset $T = \{t^1, \dots, t^r\}$ in \mathbb{TP}^{m-1} consists of $r < \infty$ points, counting copies. A sequence $(T^k, k \geq 1)$ is called a *generic perturbation* of T if each set $T^k = \{t^{k,1}, \dots, t^{k,r}\} \subset \mathbb{TP}^{m-1}$ consists of generic points and $\lim_{k \rightarrow \infty} t^{k,i} = t^i$ for all $i \in [r]$.

Lemma 3.9. *Suppose the multiset $T \subset \mathbb{TP}^{m-1}$ is finite, counting copies. Let $(T^k, k \geq 1)$ be a generic perturbation of T . Then σ is a basic cell of T if and only if there exists a sequence $(\sigma^k, k \geq 1)$, where σ^k is a full-dimensional cell in $\underline{\mathcal{H}}(-T^k)$, such that $\lim_{k \rightarrow \infty} \mathbf{d}_H(\sigma^k, \sigma) = 0$.*

In other words, the basic cells of $\underline{\mathcal{H}}(-T)$ are limits as $k \rightarrow \infty$ of full-dimensional cells in $\underline{\mathcal{H}}(-T^k)$ with respect to the Hausdorff distance. For the following proof, we shall write $\nu_T(g)$ instead of the usual shorthand $\nu(g)$ to mean the basic covector of g with respect to T .

Proof. Let r be the number of points in the multiset T , counting copies. Enumerate the points $T = \{t^1, \dots, t^r\}$. Let $T^k = \{t^{k,j}, j = 1, \dots, r\}$ be a generic perturbation from Definition 3.8. Suppose $\sigma^k \in \underline{\mathcal{H}}(-T^k)$ is a sequence of full-dimensional cells such that $\lim_{k \rightarrow \infty} \mathbf{d}_H(\sigma^k, \sigma) = 0$. Let $\nu^k := \text{coVec}_{T^k}(\sigma^k)$ be their covector. Each ν^k is a bipartite graph on $[m] \times [r]$. Since there are only finitely many such graphs, the sequence $(\nu^k, k \geq 1)$ contains a constant subsequence $\nu^{k'} = \nu'$ to which is associated the sequence $(\sigma^{k'}, k' \geq 1)$. Since each $\sigma^{k'}$ is full-dimensional, $\nu' = \nu_T(g)$ for some outcome function g . Evidently $\lim_{k' \rightarrow \infty} \mathbf{d}_H(\sigma^{k'}, \sigma) = 0$, so that σ is a basic cell of T . Conversely, suppose σ is a basic cell of T . A cell σ of $\underline{\mathcal{H}}(-T)$ has the form

$$(7) \quad \sigma = \bigcap_{j=1}^r \underline{\mathcal{H}}_{I^j}(-t^j)$$

for some subsets $I^j \subseteq [m]$, $j = 1, \dots, r$. Let g^1, \dots, g^n be all the IC outcome functions such that $\sigma = \mathcal{P}(\nu_T(g))$, for $g \in \{g^1, \dots, g^n\}$. By (7) at least one such function exists. For each k and each $i \in [n]$, consider the cell $\mathcal{P}(\nu_{T^k}(g^i))$ in $\underline{\mathcal{H}}(-T^k)$. Since T^k is generic and g^i is an outcome function, by Proposition 3.7, either this cell is empty or that it is full-dimensional. Let \mathcal{Q}^k be the set of non-empty such cells,

$$\mathcal{Q}^k = \{\mathcal{P}(\nu_{T^k}(g^i)) : \mathcal{P}(\nu_{T^k}(g^i)) \neq \emptyset, i = 1, \dots, n\}.$$

Since T^k is a generic perturbation of T , $\mathcal{Q}^k \neq \emptyset$. For each k , choose an arbitrary cell $\sigma^k \in \mathcal{Q}^k$. But $\lim_{k \rightarrow \infty} t^{k,j} = t^j$ for all $j \in [r]$, implies $\lim_{k \rightarrow \infty} \mathbf{d}_H(\sigma^k, \sigma) = 0$, so we are done. \square

The approximation result below allows us to identify the basic cells of an infinite multiset T as limit of finite approximations. For its proof it will be convenient to have another characterization of cells in $\underline{\mathcal{H}}(-T)$ using a generalization of [9, Lemma 22].




FIGURE 3. A non-generic arrangement on three types and possible generic perturbations thereof. The cells shaded gray are basic. Figure accompanies Example 3.12.

Lemma 3.10. *Introduce variables $(z_t : t \in T)$. A polyhedron σ is a cell of the tropical hyperplane arrangement $\underline{\mathcal{H}}(-T)$ with covector ν if and only if it is the projection onto the y coordinate of the set*

$$B := \{(y, z) : y_i + z_t \geq t_i \text{ for all } t \in T, i \in [m], y_i + z_t = t_i \text{ if } (i, t) \in \nu\}.$$

Proof. For each $y \in \sigma$, define $z(y)$ coordinate-wise by

$$z_t(y) = \max_{k \in [m]} \{t_k - y_k\}, \quad t \in T.$$

By definition, $y \in \sigma$ if and only if $y_i - t_i \leq y_j - t_j$ for all $(i, t) \in \nu$. For each $y \in \sigma$, the pair $(y, z(y))$ belongs to B . Conversely, if $(y, z) \in B$ for some z , then $y_i - t_i \leq y_j - t_j$, so $y \in \sigma$. \square

Proposition 3.11. *Let $T \subset \mathbb{TP}^{m-1}$ be a multiset and $\sigma \in \underline{\mathcal{H}}(-T)$. Let $(T^k, k \geq 1)$ be a sequence of finite multisets, counting copies, such that $\lim_{k \rightarrow \infty} \mathbf{d}_H(T^k, T) = 0$. Then $\sigma \in \mathbf{cells}(T)$ if and only if there is a sequence of basic cells $\sigma^k \in \mathbf{cells}(T^k)$ such that $\lim_{k \rightarrow \infty} \mathbf{d}_H(\sigma^k, \sigma) = 0$.*

Proof. Without loss of generality, one can assume that T^k is an increasing sequence, i.e. $T^k \subseteq T$, and $T^k \subseteq T^{k+1}$ for all $k \geq 1$, and that the limit is dense in T . Let σ be a cell of T with covector ν . Employing Lemma 3.10, we see that σ is a polytope in \mathbb{TP}^{m-1} , where each $t \in T$ contributes a constraint of the form

$$y_i + z_t > t_i \text{ if } (i, t) \notin \nu, \quad y_i + z_t = t_i \text{ if } (i, t) \in \nu.$$

In particular, if t, t' are arbitrarily close, then $(i, t) \in \nu$ if and only if $(i, t') \in \nu$. This implies that the covector of σ with respect to T is completely determined by its covector with respect to a dense subset of T . This implies the result. \square

We remark that by combining Proposition 3.7 and 3.11, one can choose T^k to be a finite, generic sequence that approximates T . This way, each basic cell of T is the limit of a sequence of full-dimensional cells of an approximating sequence of max-plus tropical polytopes. This is particularly useful in specific examples, as full-dimensional cells of a generic approximation of T are easy to identify. We conclude this section with some examples.

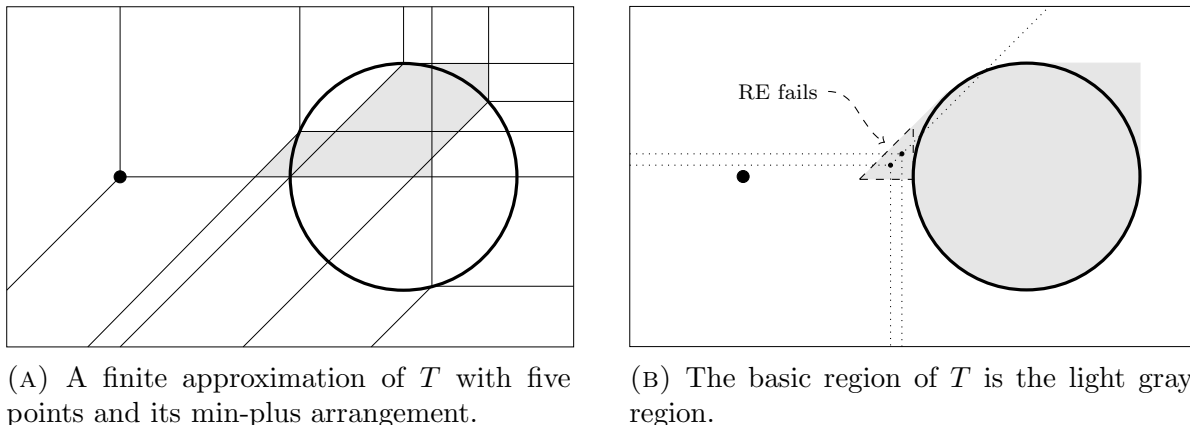


FIGURE 4. The set T , drawn in black, consists of the circle and an isolated point. Panel (A) depicts the basic cells in gray for a finite approximation of T consisting of five points. In Panel (B), the basic region is obtained by applying Proposition 3.11. Figure accompanies Example 3.14 and 4.6.

Example 3.12 (Generic perturbation of finite point-configurations). Figure 3 depicts arrangements on three points, with no repeated points. The points in Panel (A) are not generic: the projections of t_1 and t_3 onto the first coordinate is a point. Panels (B) and (C) show possible generic perturbations of these points. The covectors of the cells shaded gray are

$$\underline{\text{coVec}}_T(S) = \begin{pmatrix} 1 & 0 & 1 \\ 1 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \quad \underline{\text{coVec}}_L(S') = \begin{pmatrix} 0 & 0 & 1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \quad \underline{\text{coVec}}_L(S'') = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}$$

The cells S' and S'' are full-dimensional and thus basic by Proposition 3.7. Indeed, their covectors are outcome graphs g and h respectively, with $g(t'_1) = 2, g(t_2) = 3, g(t_3) = 1$ and $h(t_1) = 1, h(t_2) = 3, h(t'_3) = 2$. The cell S in Panel (A) is basic, as $P(\nu(g)) = P(\nu(h)) = P(\underline{\text{coVec}}_T(S))$. One can also verify via Proposition 3.11, as S can be obtained as the limit of S' as t'_1 approaches t_1 , or as the limit of S'' as t'_3 approaches t_3 .

Example 3.13 (A multiset with copies). We could also view the point configurations in Figure 3 (B) and (C) as arising from a single point with three copies that have been perturbed to become generic. In an arrangement arising from a point with three copies, the basic cell is the point itself and the covector is the 3×3 all ones matrix. Such a multiset admits $3!$ different IC mechanisms that all have the same price, namely the point we started out with.

Example 3.14 (Limiting approximation). The set $T \subset \mathbb{TP}^2$ in Figure 4 consists of the black dot and the circle drawn using opaque lines. This set T can be approximated as the closed limit of a finite sequence (T^k) . Each T^k consists of the isolated point and $k - 1$ points on the circle. Panel (A) depicts T^5 and the corresponding min-plus arrangement $\underline{\mathcal{H}}(-T^5)$. The basic cells of $\underline{\mathcal{H}}(-T^5)$ are shaded gray. In Panel (B), the basic region $\text{basic}(T)$ is shaded gray. By Proposition 3.11, it is computed by taking limits of the basic cells of T^k . Basic cells in the dotted triangular region to the left of the circle consist of parallel lines of slope $(1, 1)$. The basic cells elsewhere consist of a single point. In particular, T is not RE. In Example 4.6, we give another way to verify that T is not RE given its basic region.

4. MECHANISMS AS ALLOCATION MATRICES

We now revisit Rochet's theorem using tropical geometry. Identify an outcome function g with its allocation matrix L^g defined in (1). The main result of this section, Theorem 4.3, specifies the set of matrices that parametrizes all outcome functions on a given multiset $T \subseteq \mathbb{TP}^{m-1}$. It allows one to infer properties of g from L^g , such as incentive compatibility, the set of IC payments, and the dimension of this set, *without* computing cycle weights in L^g . This geometric insight allows us to easily construct examples of outcome functions, as illustrated in the examples below. In particular, Theorem 1.2 follows as an easy consequence.

We introduce some notational shortcuts and definitions. For a matrix $L \in \mathbb{R}^{m \times m}$ with zero diagonal, let $L_1, \dots, L_m \in \mathbb{TP}^{m-1}$ be the m rows of L , viewed as vectors in \mathbb{TP}^{m-1} . For $j, k \in [m], j \neq k$, write $\bar{\mathcal{L}}_j$ for $\bar{\mathcal{H}}_j(-L_j)$. Let $\bar{\mathcal{L}}_j^\circ$ denote the interior of this cone. Write $\bar{\mathcal{L}}_{jk}$ for $\bar{\mathcal{H}}_{jk}(-L_j)$. As before, we use the underline notation to mean the analogous quantity in min-plus. For a matrix L , define the zero eigenspace of L to be the possibly empty set

$$(8) \quad \text{Eig}_0(L) := \bigcap_{j=1}^m \underline{\mathcal{L}}_j.$$

For a multiset $T \subseteq \mathbb{TP}^{m-1}$, say that L is *realizable with respect to* T if there exists some outcome function $g : T \rightarrow [m]$ such that $L = L^g$. In that case, say that L is *realized by* g .

Definition 4.1. For $j, k \in [m], j \neq k$, define $\mathcal{I}_{jk} = \bar{\mathcal{L}}_{jk} \cap T$. A (j, k) -*witness* is a sequence $\{s^{j,r} : r \geq 1\} \subseteq T \cap \bar{\mathcal{L}}_j^\circ$ such that

$$\lim_{r \rightarrow \infty} d(s^{j,r}, \bar{\mathcal{L}}_{jk}) = 0.$$

We say that L *separates* T at (j, k) if

$$(9) \quad d(T \cap \bar{\mathcal{L}}_j, \bar{\mathcal{L}}_{jk}) = 0,$$

and in addition, whenever $\mathcal{I}_{jk} = \mathcal{I}_{kj} = \{s\}$ for some $s \in \mathbb{TP}^{m-1}$, then there exists a (j, k) -witness or a (k, j) -witness. Say that L *separates* T if L separates T for all $j, k \in [m], j \neq k$.

Definition 4.2. For $p \in \mathbb{TP}^{m-1}$, the graph of p (with respect to T) is the directed graph on m nodes, with edge (i, j) if and only if

$$d(T \cap \bar{\mathcal{L}}_i, \bar{\mathcal{H}}_{ij}(-p)) = 0.$$

Theorem 4.3. Let $L \in \mathbb{R}^{m \times m}$ be a matrix with zero diagonal, $T \subseteq \mathbb{TP}^{m-1}$ be a multiset.

(1) L is realizable if and only if L separates T , and

$$T \subseteq \bigcup_{k=1}^m \bar{\mathcal{L}}_k.$$

(2) Suppose L is realized by g . Then g is weakly monotone if and only if the open sets $\bar{\mathcal{L}}_1^\circ, \dots, \bar{\mathcal{L}}_m^\circ$ are pairwise disjoint.

(3) Suppose L is realized by g . The set of incentive compatible payments of g is $\text{Eig}_0(L)$. In particular, g is IC if and only if $\text{Eig}_0(L) \neq \emptyset$.

(4) Suppose L is realized by g and $\text{Eig}_0(L) \neq \emptyset$. Let p be a point in the relative interior of $\text{Eig}_0(L)$. Then the dimension of $\text{Eig}_0(L)$ is the number of strongly connected components in the graph of p with respect to T minus 1.

We defer the proof to Section 4.1. Instead, we illustrate the witnessing condition and the implications of the theorem with examples. They show that it is simple to construct mechanisms and verify realizability, incentive compatibility and revenue equivalence.

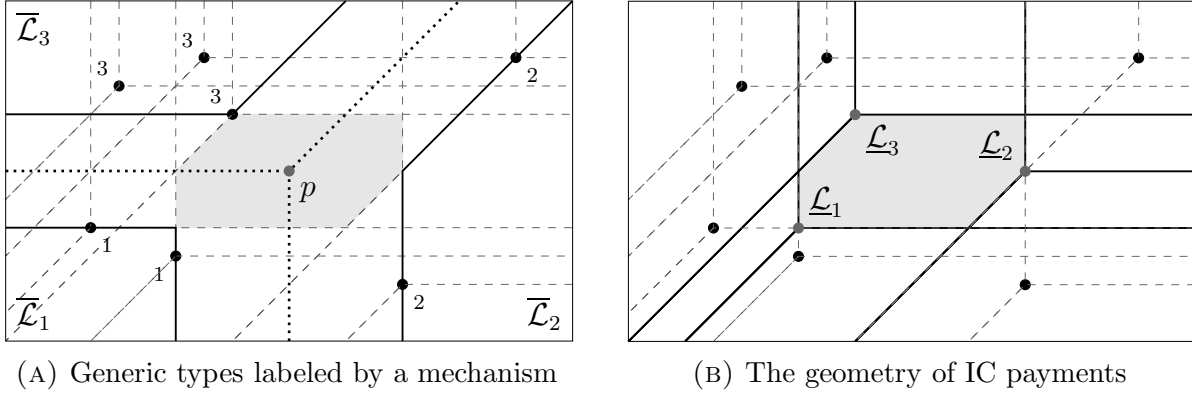


FIGURE 5. Geometric construction of allocation matrices. The type set consists of the black dots that are apices of the min-plus hyperplanes drawn using dashed lines. In Panel (A), types have been labeled according to their outcome under an outcome function g . In Panel (B), the gray dots are the rows of the matrix L^g , viewed as points in \mathbb{TP}^2 . The set of IC payments of g is shaded gray in both panels. Figure accompanies Example 4.4.

Example 4.4 (Geometric construction of allocation matrices). Figure 5 (A) shows how to construct the apices of the sectors $\bar{\mathcal{L}}_i$ geometrically. The set T consists of the black generic points. The dashed lines are hyperplanes of $\mathcal{H}(-T)$. The labels next to the points define an outcome function $g : T \rightarrow [3]$. For $i \in [3]$, the heavy black lines define the boundary of the max-plus sectors $\bar{\mathcal{L}}_i$, whose apex is the i -th row of L^g . For example, for $i = 1$, the first row of L^g can be written as $L_1^g = (0, \sup_{t \in g^{-1}(1)} \{t_2 - t_1\}, \sup_{t \in g^{-1}(1)} \{t_3 - t_1\})$. The set of IC payments of g equals the basic cell shaded gray, computed using Theorem 1.1. Panel (B) verifies that this basic cell equals $\text{Eig}_0(L)$, as stipulated by Theorem 4.3, part (2). Here $\text{Eig}_0(L)$ is a polytope in \mathbb{TP}^2 of dimension 2. To verify with Theorem 4.3 part (4), let p be a point in the interior of $\text{Eig}_0(L)$ as shown in Panel (A), with $\bar{\mathcal{H}}(-p)$ shown in dotted lines. One can readily verify that the graph of p consists of three points with no edges, thus it has three strongly connected components. So $\text{Eig}_0(L)$ has dimension 2, as expected.

Example 4.5 (Realizable and non-realizable matrices). In Figure 6, $T = \{t_1, t_2, t_3, t_4\}$ is the set of four black points. Each panel defines a matrix L whose rows give rise to the sectors $\bar{\mathcal{L}}_1, \bar{\mathcal{L}}_2$ and $\bar{\mathcal{L}}_3$. The matrix in Panel (A) is realizable: it equals L^g for the mechanism g is defined by $g(t_2) = g(t_3) = 3$, $g(t_1) = 1$ and $g(t_4) = 2$. The matrix L in Panel (B) is not realizable. Suppose for contradiction that it is realized by some outcome function g . By Theorem 4.3, we must have $g(t_4) = 2$, $g(t_1) = 1$, $g(t_3) = 3$, and $g(t_2)$ is either 1 or 3. If $g(t_2) = 3$, then L^g must equal that shown in Figure 6 (A) and thus is not equal to L . So $g(t_2) = 1$, but computations show that for this g , $L^g \neq L$. Thus L is not realizable. Realizability fails since L does not separate T at $\{1, 3\}$. Here, one has $\mathcal{I}_{13} = \bar{\mathcal{L}}_{13} \cap \bar{\mathcal{L}}_{31} \cap T = \{t_2\}$, but there is no $(1, 3)$ - nor a $(3, 1)$ -witness.





FIGURE 6. Realizable and non-realizable matrices, defined by the apices of $\overline{\mathcal{L}}_i$ for $i \in [3]$. Figure accompanies Example 4.5.

Example 4.6 (A type space without RE). Consider Figure 4 (B), where the basic region of a certain type set $T \subset \mathbb{TP}^2$ is shaded in gray. To each point p in this basic region, put a max-plus hyperplane and compute its graph. We find that for all p in the circle, the graph of p is the complete graph on 3 nodes, and for p in the top right region outside the circle, the graph has edges $(1,3), (3,1), (2,3)$ and $(3,2)$. These graphs are all strongly connected, thus for each such p , outcome functions with IC payment p is RE. On the other hand, for p in the dotted triangular region to the left of the circle, the graph has edges $(2,3)$ and $(3,2)$. In Figure 4 (B), we show two such p as black dots, their max-plus hyperplanes in dotted lines. By Theorem 1.1 part (1) and 4.3 part (4), the set of IC payments for outcome functions with IC payment p has dimension 1. In particular, such outcome functions are not RE. So T is not RE. One can also verify the dimension of the basic cells via finite approximation, as done in Example 3.14.



(A) A mechanism that is weakly monotone but not IC. (B) The min-plus arrangement verifying non-IC. (C) On the same type space, another mechanism that is RE.

FIGURE 7. Demonstration of weak monotonicity, lack of incentive compatibility, on a revenue equivalent domain. Figure accompanies Example 4.7.

Example 4.7 (Weak monotonicity, IC and an RE type space). In Figure 7, the set T consists of all points in the region shaded gray. This type space T is RE, but not all weakly monotone

outcome functions on T are IC, on convex type spaces this is cannot happen *cf.* [23]. Panel (A) defines a weakly monotone outcome function via L^g . Weak monotonicity follows from part (2) of Theorem 4.3, since the open sectors $\overline{\mathcal{L}}_1^\circ, \overline{\mathcal{L}}_2^\circ, \overline{\mathcal{L}}_3^\circ$ are pairwise disjoint. Panel (B) verifies that $\text{Eig}_0(L^g) = \emptyset$, so g is not IC, by Theorem 4.3 part (3). The basic region of T is all of \mathbb{TP}^2 . By considerations of various max-plus hyperplanes, such as the one depicted in Figure 7(C), one can confirm that T is RE. Note that the closure of this type space is not path-connected, nor ‘boundedly grid-wise connected’, which are other known and easily verified sufficient conditions for a type space to be RE, *cf.* [4, Theorems 1 and 4].

4.1. Proofs of Theorem 4.3 and 1.2.

Proof of Theorem 4.3, part (1). Suppose L is realized by g . For each $j = 1, 2, \dots, m$, $g^{-1}(j) \subseteq \overline{\mathcal{L}}_j$, since $\max_{k \in [m]}(L_{jk} + t_k) = L_j \odot t = t_j$. The sets $g^{-1}(1), \dots, g^{-1}(m)$ partition T , so $T \subseteq \bigcup_{j=1}^m \overline{\mathcal{L}}_j$. It remains to show that L separates T at an arbitrary pair $\{j, k\}$, where $j, k \in [m], j \neq k$. We have that $d(T \cap \overline{\mathcal{L}}_j, \overline{\mathcal{L}}_{jk}) = 0$ since $L_{jk} = \inf_{t \in g^{-1}(j)} \{t_j - t_k\}$. If $\mathcal{I}_{jk} = \mathcal{I}_{kj} = \{s\}$, then if $g(s) = j$, the infimum in the definition of L_{kj} must be achieved by a (k, j) -witness, so there must exist a (k, j) -witness. Conversely, if $g(s) = k$, then there must exist a (j, k) -witness. This shows that L separates T at $\{j, k\}$, as desired. For the converse direction, suppose L is a matrix with zero diagonal such that $T \subseteq \bigcup_{j=1}^m \overline{\mathcal{L}}_j$ and L separates T . Define $g : T \rightarrow [m]$ as follows. For a point $t \in T \cap \overline{\mathcal{L}}_j$, let $g(t) = j$. The remaining points must lie on $\bigcup_{j,k \in [m], j \neq k} \mathcal{I}_{jk}$ by definition of \mathcal{I}_{jk} ’s. Assign these points such that points on \mathcal{I}_{jk} have either outcome j or k , and such that on every non-empty boundary \mathcal{I}_{jk} there exists a point with outcome j . The only case where this cannot be done is if $\mathcal{I}_{jk} = \mathcal{I}_{kj} = \{s\}$. In this case, if there is a (j, k) -witness, set $g(s) = k$, else, since L separates T , there must exist a (k, j) -witness, so set $g(s) = j$. We claim that L is realized by g . Fix $j, k \in [m], j \neq k$. By definition of $g^{-1}(j)$,

$$(10) \quad \inf_{s \in g^{-1}(j)} (s_j - s_k) \geq L_{jk}.$$

Furthermore, there must be a point in \mathcal{I}_{jk} or a (j, k) -witness in $g^{-1}(j)$. So (10) holds with an equality, thus $L = L^g$, as claimed. \square

Proof of Theorem 4.3, part (2). Fix a pair of indices $j, k \in [m], j \neq k$. We claim that

$$\begin{aligned} L_{jk} + L_{kj} &> 0 \Leftrightarrow \overline{\mathcal{L}}_j \cap \overline{\mathcal{L}}_k = \emptyset, \\ L_{jk} + L_{kj} &= 0 \Leftrightarrow \overline{\mathcal{L}}_j \cap \overline{\mathcal{L}}_k \text{ on their boundaries, i.e. } \overline{\mathcal{L}}_j \cap \overline{\mathcal{L}}_k = \partial \overline{\mathcal{L}}_j \cap \partial \overline{\mathcal{L}}_k \\ L_{jk} + L_{kj} &< 0 \Leftrightarrow \overline{\mathcal{L}}_j \cap \overline{\mathcal{L}}_k \text{ in their interiors, i.e. } \overline{\mathcal{L}}_j^\circ \cap \overline{\mathcal{L}}_k^\circ \neq \emptyset. \end{aligned}$$

This claim implies statement (4) by definition of weakly monotone. Now let us prove the claim. Note that $\overline{\mathcal{L}}_j$ and $\overline{\mathcal{L}}_k$ are closed polyhedra, so they either do not intersect, intersect on their boundaries, or intersect in their interiors. So it is sufficient to prove the last two equivalences in the statements. Suppose there exists $t \in \overline{\mathcal{L}}_j \cap \overline{\mathcal{L}}_k$. Then by definition of the sectors,

$$L_{jk} + L_{kj} \leq (t_j - t_k) + (t_k - t_j) = 0.$$

In particular, strict equality holds if and only if t lies on the boundary of $\overline{\mathcal{L}}_j \cap \overline{\mathcal{L}}_k$, while strict inequality holds if and only if t lies in either of their interiors. In that case, one can pick a $t' \in \overline{\mathcal{L}}_j^\circ \cap \overline{\mathcal{L}}_k^\circ$, then $L_{jk} + L_{kj} < (t'_j - t'_k) + (t'_k - t'_j) = 0$. This proves the desired statement. \square

Proof of Theorem 4.3, part (3) and (4). By definition,

$$\text{Eig}_0(L) = \{p \in \mathbb{TP}^{m-1} : L \odot p = p\}.$$

Part (3) of the theorem follows from uniqueness of the tropical eigenvalue [8]. For part (4), by Proposition 3.3, $\text{Eig}_0(L^g) = \mathcal{P}(\nu(g))$, where $\nu(g) = \text{coVec}_T(p)$. Note that

$$(11) \quad d(\overline{\mathcal{H}}_{ij}(-p), T \cap \overline{\mathcal{L}}_i) = \left(\inf_{t \in g^{-1}(i)} (t_i - t_j) \right) - (p_i - p_j) = L_{ij}^g - (p_i - p_j).$$

Therefore, the graph of p has edge (i, j) if and only if $p_i - p_j = L_{ij}^g$. Form the matrix L' via $L'_{ij} = L_{ij}^g - (p_i - p_j)$. Then L' has eigenvalue 0, and any cycle in the graph of p is a zero-length cycle. In tropical linear systems terminology, the graph of p is the critical graph of L' . It follows from [5, Theorem 4.3.3, 4.3.5] that the dimension of the eigenspace of L' equals the number of connected components of L' . But $\text{Eig}(L') = -p + \text{Eig}(L^g)$, so $\text{Eig}(L^g)$ has the same dimension. This proves the theorem. \square

Proof of Theorem 1.2. This follows immediately from Theorem 4.3, part (4). If T is RE, then for every $p \in \text{basic}(T)$, its graph equals the graph of L^g for some RE mechanism (g, p) , which is connected. Conversely, suppose the graph of every p in $\text{basic}(T)$ is strongly connected. Take a IC mechanism (g, p) , with p in the relative interior of $\mathcal{P}(g)$. As its graph is connected, $\text{Eig}_0(L^g) = \{p\}$, so (g, p) is RE. \square

5. ACKNOWLEDGEMENTS

We wish to thank John Weymark for suggesting to explore possible connections of mechanism design to tropical geometry. Our geometric approach is inspired by his work with coauthors in [7]. Part of this work was written during the Hausdorff Center Conference “Tropical Geometry and Economics”. We thank the participants for helpful discussions. The Hausdorff Center’s role in catalyzing research in mathematical economics is gratefully acknowledged. Finally we thank Paul Edelman and John Weymark for pointing out an error in an earlier version, and Alexey Kushnir for generously sharing his manuscript.

REFERENCES

- [1] Marianne Akian, Stephane Gaubert, and Alexander Guterman. Tropical polyhedra are equivalent to mean payoff games. *International Journal of Algebra and Computation*, 22(01):1250001, 2012. 3
- [2] François Baccelli, Guy Cohen, Geert Jan Olsder, and Jean-Pierre Quadrat. *Synchronization and linearity: an algebra for discrete event systems*. John Wiley & Sons Ltd, 1992. 3
- [3] Elizabeth Baldwin and Paul Klemperer. Understanding preferences: “demand types”, and the existence of equilibrium with indivisibilities. Technical report, The London School of Economics and Political Science, 2016. 3
- [4] Sushil Bikhchandani, Shurojit Chatterji, Ron Lavi, Ahuva Mu’alem, Noam Nisan, and Arunava Sen. Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica*, 74(4):1109–1132, 2006. 16
- [5] Peter Butkovič. *Max-linear Systems: Theory and Algorithms*. Springer Monographs in Mathematics, 2012. 3, 17
- [6] Kim-Sau Chung and Wojciech Olszewski. A non-differentiable approach to revenue equivalence. *Theoretical Economics*, 2(4):469–487, 2007. 2
- [7] Katherine Cuff, Sunghoon Hong, Jesse A. Schwartz, Quan Wen, and John A. Weymark. Dominant strategy implementation with a convex product space of valuations. *Social Choice and Welfare*, 39(2):567–597, 2012. 17
- [8] Raymond A. Cuninghame-Green. Describing industrial processes with interference and approximating their steady-state behaviour. *OR*, 13(1):95–100, 1962. 4, 5, 17

- [9] Mike Develin and Bernd Sturmfels. Tropical convexity. *Doc. Math.*, 9:1–27, 2004. 6, 10
- [10] Alex Fink and Felipe Rincón. Stiefel tropical linear spaces. *Journal of Combinatorial Theory, Series A*, 135:291 – 331, 2015. 6
- [11] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991. 1
- [12] Stéphane Gaubert and Ricardo D. Katz. Minimal half-spaces and external representation of tropical polyhedra. *Journal of Algebraic Combinatorics*, 33(3):325–348, 2011. 6
- [13] Alex Gershkov, Benny Moldovanu, and Xianwen Shi. Optimal voting rules. Forthcoming in *Review of Economic Studies*. 3
- [14] Michael Joswig. The cayley trick for tropical hypersurfaces with a view toward ricardian economics. arXiv:1606.09165. 3
- [15] Michael Joswig. *Essentials of Tropical Combinatorics*. in preparation. 3, 6
- [16] Michael Joswig and Georg Loho. Weighted digraphs and tropical cones. *Linear Algebra and its Applications*, 501:304 – 343, 2016. 6
- [17] Jean-Jacques Laffont and David Martimort. *The theory of incentives: the principal-agent model*. Princeton university press, 2009. 3
- [18] Diane Maclagan and Bernd Sturmfels. *Introduction to Tropical Geometry*. Graduate Studies in Mathematics. American Mathematical Society, 2015. 3, 6, 10
- [19] Hervé Moulin. On strategy-proofness and single peakedness. *Public Choice*, 35(4):437–455, 1980. 3
- [20] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981. 1, 3
- [21] Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of economic theory*, 29(2):265–281, 1983. 3
- [22] Jean-Charles Rochet. A necessary and sufficient condition for rationalizability in a quasilinear context. *J. Math. Econom.*, 16(2):191–200, 1987. 1, 5
- [23] Michael Saks and Lan Yu. Weak monotonicity suffices for truthfulness on convex domains. In *Proceedings of the 6th ACM Conference on Electronic Commerce*, EC '05, pages 286–293. ACM, 2005. 16
- [24] Yoshinori Shiozawa. International trade theory and exotic algebras. *Evolutionary and Institutional Economics Review*, 12(1):177–212, 2015.
- [25] Ngoc M. Tran and Josephine Yu. Product-mix auctions and tropical geometry. 2015. arXiv:1505.05737. 3
- [26] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.
- [27] Rakesh V. Vohra. *Mechanism Design: A Linear Programming Approach*. Econometric Society Monographs. Cambridge University Press, 2011. 1, 3, 5

(Robert Alexander Crowell) HAUSDORFF CENTER FOR MATHEMATICS, BONN 53115, GERMANY

E-mail address: crowellr@uni-bonn.de

(Ngoc Mai Tran) DEPARTMENT OF MATHEMATICS, UNIVERSITY OF TEXAS AT AUSTIN, TEXAS TX 78712, USA AND HAUSDORFF CENTER FOR MATHEMATICS, BONN 53115, GERMANY

E-mail address: ntran@math.utexas.edu